

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 September 2008 (04.09.2008)

PCT

(10) International Publication Number
WO 2008/104549 A2

(51) International Patent Classification: Not classified

(21) International Application Number:
PCT/EP2008/052329

(22) International Filing Date:
26 February 2008 (26.02.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/892,238 28 February 2007 (28.02.2007) US

(71) Applicant (for all designated States except US): **FOTO-
NATION VISION LIMITED** [IE/IE]; Galway Business
Park, Dangan, Galway (IE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **IONITA, Mircea**
[RO/IE]; 75 Woodfield, Knocknacarra, Galway (IE).
BACIVAROV, Ioana [RO/IE]; 75 Rosan Glas, Rahoon,
Galway (IE). **CORCORAN, Peter** [IE/IE]; Cregg Clare-
galway, Galway (IE).

(74) Agents: **BOYCE, Conor** et al.; c/o F. R. Kelly & Co., 27
Clyde Road, Dublin 4 (IE).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA,
CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE,
EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID,
IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC,
LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN,
MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH,
PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV,
SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN,
ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL,
NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG,
CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declaration under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a
patent (Rule 4.17(ii))

Published:

— without international search report and to be republished
upon receipt of that report

(54) Title: SEPARATING DIRECTIONAL LIGHTING VARIABILITY IN STATISTICAL FACE MODELLING BASED ON
TEXTURE SPACE DECOMPOSITION

(57) Abstract: A technique for determining a characteristic of a face or certain other object within a scene captured in a digital image including acquiring an image and applying a linear texture model that is constructed based on a training data set and that includes a class of objects including a first subset of model components that exhibit a dependency on directional lighting variations and a second subset of model components which are independent of directional lighting variations. A fit of the model to the face or certain other object is obtained including adjusting one or more individual values of one or more of the model components of the linear texture model. Based on the obtained fit of the model to the face or certain other object in the scene, a characteristic of the face or certain other object is determined.



WO 2008/104549 A2

Separating Directional Lighting Variability In Statistical Face Modelling Based On Texture
Space Decomposition

5

BACKGROUND

The appearance of an object can be represented by statistical models trained using a set of annotated image examples. This is thus highly dependent on the way in which the model is trained. A new image can be interpreted by finding the best plausible match of the model to the image data. While there has been a great deal of literature in computer vision detailing methods for handling statistical models for human faces, there still exist some problems wherein solutions are desired. For example, statistical models for human faces are sensitive to illumination changes, especially if lighting in the test image differs significantly from conditions learned from a training set. The appearance of a face can change dramatically as lighting conditions change. Due to the 3D aspect of the face, a direct lighting source can cast strong shadows and shading which affect certain facial features. Variations due to illumination changes can be even greater than variations between the faces of two different individuals.

Various methods have been proposed to overcome this challenge. A feature-based approach seeks to utilize features that are invariant to lighting variations. In C. Hu, R. Feris, and M. Turk, "Active wavelet networks for face alignment," in *Proc. of the British Machine Vision Conference*, East Eaglia, Norwich, UK, 2003, it is proposed to replace the AAM texture by an active wavelet network for face alignment, while in S. Le Gallou, G. Breton, C. Garcia, and R. S'eguier, "Distance maps: A robust illumination preprocessing for active appearance models," in *VISAPP'06, First International Conference on Computer Vision Theory and Applications*, Set'ubal, Portugal, 2006, vol. 2, pp. 35–40, texture is replaced by distance maps that are robust against lighting variations.

Other methods rely on removing illumination components using lighting models. The linear subspace approaches of S.Z. Li, R. Xiao, Z.Y. Li, and H.J. Zhang, "Nonlinear mapping of multi-view face patterns to a Gaussian distribution in a low dimensional space," in *RATFG-*

RTS '01: Proceedings of the IEEE ICCV Workshop on Recognition, Analysis, and Tracking of Faces and Gestures in Real-Time Systems, 2001, p. 47, and M. Bichsel, "Illumination invariant object recognition," in *ICIP '95: Proceedings of the 1995 International Conference on Image Processing - Vol. 3*, 1995, p. 3620, and P.N. Belhumeur, J. Hespanha, and D.J. Kriegman, "Eigenfaces vs. fisherfaces: Recognition using class specific linear projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, 1997, approximate the human face surface with a Lambertian surface and compute a basis for a 3D illumination subspace, using images acquired under different lighting conditions.

The illumination convex cone goes a step further with the model, taking into account shadows and multiple lighting sources, as in P.N. Belhumeur and D.J. Kriegman, "What is the set of images of an object under all possible lighting conditions?," in *CVPR '96: Proceedings of the 1996 Conference on Computer Vision and Pattern Recognition*, 1996, p. 270, and A.S. Georghiades, D.J. Kriegman, and P.N. Belhumeur, "Illumination cones for recognition under variable lighting: Faces," in *CVPR '98: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1998, p. 52, and A.S. Georghiades, P.N. Belhumeur, and D.J. Kriegman, "From few to many: Generative models for recognition under variable pose and illumination.," in *FG*, 2000, pp. 277–284.

More complex models have been proposed like the geodesic illumination basis model of R. Ishiyama and S. Sakamoto, "Geodesic illumination basis: Compensating for illumination variations in any pose for face recognition.," in *ICPR (4)*, 2002, pp. 297–301, or the 3D linear subspace model that segments the images into regions with directions of surface normals close to each other as in A.U. Batur and M.H. Hayes, "Linear subspaces for illumination robust face recognition.," in *CVPR (2)*, 2001, pp. 296–301.

The canonical form approach appears as an alternative, where an attempt to normalize variations in appearance by image transformations or by synthesizing a new image from the given image in a normalized form is undertaken. Recognition is then performed using this canonical form as in W. Zhao, *Robust image based 3d face recognition*, Ph.D. thesis, 1999, Chair-Rama Chellappa or W. Gao, S. Shan, X. Chai, and X. Fu, "Virtual face image generation for illumination and pose insensitive face recognition," *ICME*, vol. 3, pp. 149–

152, 2003.

In T. Shakunaga and K. Shigenari, "Decomposed eigenface for face recognition under various lighting conditions," *CVPR*, vol. 01, pp. 864, 2001, and T. Shakunaga, F. Sakaue, and
5 K. Shigenari, "Robust face recognition by combining projection-based image correction and decomposed eigenface," 2004, pp. 241–247, decomposition of an eigenface into two orthogonal eigenspaces is proposed for realizing a general face recognition technique, under lighting changes. A somewhat similar approach is used in J.M. Buenaposada, E. Munoz, and
10 L. Baumela, "Efficiently estimating facial expression and illumination in appearance-based tracking," 2006, p. I:57, for face tracking, where the face is represented by the addition of two approximately independent subspaces to describe facial expressions and illumination, respectively.

In N. Costen, T.F. Cootes, G.J. Edwards, and C.J. Taylor, "Automatic extraction of the face
15 identity-subspace.," in *BMVC*, 1999, and N. Costen, T.F. Cootes, G.J. Edwards, and C.J. Taylor, "Simultaneous extraction of functional face subspaces," *CVPR*, vol. 01, pp. 1492, 1999, facial appearance models of shape and texture are employed and non-orthogonal texture subspaces for lighting, pose, identity, and expression are extracted using appropriate image sets. An iterative expectation-maximization algorithm is then applied in order to
20 maximize the efficiency of facial representation over the added subspaces. The projections on each subspace are then used to recalculate the subspaces. This approach is shown to improve the identity recognition results. It is still desired to have an algorithm that permits less complex handling of illumination changes, and obtaining a general and robust facial appearance model.

25

PCA-based models generally do not decouple different types of variations. AAM techniques are using PCA, and thus inherit this limitation of being practically incapable of differentiating among various causes of face variability, both in shape and texture. An important drawback of a non-decoupled PCA-based model is that it can introduce non-valid space regions,
30 allowing the generation of non-realistic shape/texture configurations. Moreover, the interpretation of the parameters of the global model can be ambiguous, as there is no clear distinction of the kind of variation they stand for. It is recognized by the inventors that it would be desirable to obtain specialized subspaces, such as an identity subspace and/or a

directional lighting subspace.

Changes in lighting or illumination represent one of the most complex and difficult to analyze sources of face variability. Thus it is desired to decouple variations in identity from those caused by directional lighting. It is further desired to split the shape model by decoupling identity from pose or expression. It is recognized by the inventors that decoupling the pose variations from the global shape model can be realized by using a proper training set, in which the individuals are presented in several poses, normally covering a range within 30° - 40° for head tilting.

SUMMARY OF THE INVENTION

A technique is provided for determining a characteristic of a face or certain other object within a scene captured in a digital image. A digital image is acquired including a face or certain other object within a scene. A linear texture model is applied that is constructed based on a training data set and that includes a class of objects including a first subset of model components that exhibit a dependency on directional lighting variations and a second subset of model components which are independent of directional lighting variations. An initial location of the face or certain other object in the scene is initially determined. A fit of the model to the face or certain other object is obtained as one or more individual values of one or more of the model components of the linear texture model is/are adjusted. Based on the obtained fit of the model to the face or certain other object in the scene, at least one characteristic of the face or certain other object is determined. A corrected image including the determined characteristic is electronically stored, transmitted, has a face or other object recognition program applied thereto, edited, or displayed, or combinations thereof.

A further technique is provided for adjusting a characteristic of a face or certain other object within a scene captured in a digital image is also provided. A digital image is acquired including a face or certain other object within a scene. A linear texture model is obtained that is constructed based on a training data set and includes a class of objects including a first subset of model components that exhibit a dependency on directional lighting variations and a second subset of model components which are independent of directional lighting variations. An initial location of the face or certain other object in the scene is determined. A

fit of the model to the face or certain other object in the scene is obtained as one or more individual values of one or more model components of the linear texture model is/are adjusted. Based on the obtained fit of the model to the face or certain other object in the scene, a characteristic of the face or certain other object is adjusted as one or more values of one or more model components of the linear texture model are changed to generate an adjusted object model. The adjusted object model is superimposed onto the digital image. The method includes electronically storing, transmitting, applying a face recognition program to, editing, or displaying the corrected face image, or combinations thereof.

- 10 The model components may include eigenvectors, and the individual values may include eigenvalues of the eigenvectors. The determined characteristic may include a feature that is independent of directional lighting. A reconstructed image without a periodic noise component may be generated.
- 15 A second fit may be obtained of the face or certain other object to a second linear texture model that is based on a training data set and that includes a class of objects including a set of model components which lack a periodic noise component. The method may include extracting the periodic noise component including determining a difference between the face or certain other object and the reconstructed image. A frequency of the noise component may be determined, and the periodic noise component of the determined frequency is removed.

- 25 An exposure value for the face or certain other object may be determined, as a fit is obtained of the face or certain other object to a second linear texture model that is based on a training data set and that includes a class of objects including a set of model components that exhibit a dependency on exposure value variations. An effect of a background region or density contrast caused by shadow, or both, may be reduced.

- 30 The method may include controlling a flash to accurately reflect a lighting condition. A flash control condition may be obtained by referring to a reference table. A flash light emission may be controlled according to the flash control condition. An effect of contrasting density caused by shadow or black compression or white compression or combinations thereof may be reduced.

The method may further include adjusting or determining a sharpness value, or both. A second linear texture model may be obtained that is constructed based on a training data set. The model may include a class of objects including a subset of model components that exhibit a dependency on sharpness variations. A fit of said second model to the face or certain other object in the scene may be obtained and one or more individual values of one or more model components of the second linear texture model may be adjusted. Based on the obtained fit of the second model to the face or certain other object in the scene, a sharpness of the face or certain other object may be adjusted as one or more values of one or more model components of the second linear texture model are changed to generate a further adjusted object model.

The method may include removing a blemish from a face or certain other object. A second linear texture model may be obtained that is constructed based on a training data set and includes a class of objects including a subset of model components that do not include the blemish. A fit of the second model to the face or certain other object in the scene may be obtained as one or more individual values of one or more model components of the second linear texture model are adjusted. Based on the obtained fit of the second model to the face or certain other object in the scene, the blemish is removed from the face or certain other object as one or more values of one or more model components of the second linear texture model are changed to generate a further adjusted object model. The blemish may include an acne blemish or other skin blemish or a photographic artefact.

A graininess value may be adjusted and/or determined. A second linear texture model may be obtained that is constructed based on a training data set and includes a class of objects including a subset of model components that exhibit a dependency on graininess variations. A fit of the second model to the face or certain other object in the scene may be obtained including adjusting one or more individual values of one or more model components of the second linear texture model. Based on the obtained fit of the second model to the face or certain other object in the scene, a graininess of the face or certain other object may be adjusted as one or more values of one or more model components of the second linear texture model are changed to generate a further adjusted object model.

A resolution value may be converted, adjusted and/or determined. A second linear texture

model may be obtained that is constructed based on a training data set and includes a class of objects including a subset of model components that exhibit approximately a same resolution as the face or certain other object. A fit of the second model to the face or certain other object in the scene may be obtained as one or more individual values of one or more model components of the second linear texture model are adjusted. Based on the obtained fit of the second model to the face or certain other object in the scene, a resolution of the face or certain other object may be converted as one or more values of one or more model components of the second linear texture model are changed to generate a further adjusted object model.

In the second technique, the adjusting may include changing one or more values of one or more model components of the first subset of model components to a set of mean values, and thereby adjusting directional lighting effects on the scene within the digital image. The first technique may also include the adjusting. The adjusting of directional lighting effects may include increasing or decreasing one or more directional lighting effects. The adjusting may include filtering directional light effects to generate a directional light filtered face image, and a face recognition program may be applied to the filtered face image.

A further method is provided for constructing a linear texture model of a class of objects, including a first subset of model components that exhibit a dependency on directional lighting variations and a second subset of model components that are independent of directional lighting variations. A training set is provided including multiple object images wherein various instances of each object cover a range of directional lighting conditions. The method also includes applying to the images a linear texture model constructed from object images each captured under uniform lighting conditions and forming a uniform lighting subspace (ULS). A set of residual texture components are determined between object images captured under directional lighting conditions and the linear texture model constructed from object images each captured under uniform lighting conditions. An orthogonal texture subspace is determined from residual texture components to form a directional lighting subspace (DLS). The uniform lighting subspace (ULS) is combined with the directional lighting subspace (DLS) to form a new linear texture model.

The method may further include:

(i) applying the new linear texture model that is constructed based on a training data set and comprises a class of objects including a subset of model components that has uniform lighting subspace (ULS) and directional lighting subspace (DLS) components,

(ii) obtaining a fit of said new model to the face or certain other object in the scene
5 including adjusting one or more individual values of one or more model components of the new linear texture model;

(iii) based on the obtained fit of the new model to the face or certain other object in the scene, determining a characteristic of the face or certain other object, and

(iv) electronically storing, transmitting, applying a face or other object recognition
10 program to, editing, or displaying the corrected face image or certain other object including the determined characteristic, or combinations thereof.

The method may further include:

(v) further comprising changing one or more values of one or more model
15 components of the new linear texture model to generate a further adjusted object model;

(vi) obtaining a fit of the new model to the face or certain other object in the scene including adjusting one or more individual values of one or more model components of the new linear texture model; and

(vii) based on the obtained fit of the new model to the face or certain other object in
20 the scene, determining a characteristic of the face or certain other object.

The model components may include eigenvectors, and the individual values may include eigenvalues of the eigenvectors.

25 A face illumination normalization method is also provided. A digital image is acquired including data corresponding to a face that appears to be illuminated unevenly. Separate sets of directional and uniform illumination classifier programs are applied to the face data. The face data are identified as corresponding to a projection of the face within the digital image on one or a combination of the directional illumination classifier programs plus a constant
30 vector representing the face according to one or a combination of the uniform illumination classifier programs, thereby decomposing the face data into orthogonal subspaces for directional and uniform illumination. An illumination condition may be normalized for the face including setting one or more illumination parameters of the directional illumination

projection to zero. The method may further include electronically storing, transmitting, applying a face recognition program to, editing, or displaying the corrected face image, or combinations thereof.

- 5 The applying may include projecting the face data onto the uniform lighting classifier program set, and then applying residual data of the face data to the directional lighting classifier program set. A face recognition program may be applied to the normalized face image. A set of feature detector programs may be applied to reject non-face data from being identified as face data. An illumination condition may be determined based on acceptance of
10 the face data by one or a combination of the directional illumination classifier programs.

The digital image may be one of multiple images in a series that include the face. The normalizing may be applied to a different image in the series than the original digital image within which the face data is identified.

15

- A face detection method is also provided. A digital image is acquired, and a sub-window is extracted from the image. Separate sets of two or more shortened face detection classifier cascades are applied. One set is trained to be selectively sensitive to a characteristic of a face region, and another set of face detection classifier cascades are insensitive to the
20 characteristic. The face data are identified as corresponding to a projection of the face within the digital image on one or a combination of the characteristic-sensitive classifier cascades plus a constant vector representing the face according to one or a combination of the characteristic-insensitive classifier cascades, thereby decomposing the face data into orthogonal subspaces for characteristic-sensitive and characteristic-insensitive conditions.
25 Based on the applying and identifying, a probability is determined that a face with a certain form of the characteristic is present within the sub-window. Based on the determining, an extended face detection classifier cascade trained for sensitivity to the form of said characteristic is applied. A final determination that a face exists within the image sub-window is provided. The process is repeated for one or more further sub-windows from the
30 image or one or more further characteristics, or both.

The characteristic or characteristics may include a directional illumination of the face region, an in-plane rotation of the face region, a 3D pose variation of the face region, a degree of

smile, a degree of eye-blinking, a degree of eye-winking, a degree of mouth opening, facial blurring, eye-defect, facial shadowing, facial occlusion, facial color, or facial shape, or combinations thereof.

5 The characteristic may include a directional illumination, and the method may include determining an uneven illumination condition by applying one or more uneven illumination classifier cascades. A front illumination classifier cascade may be applied. An illumination condition of a face within a sub-window may be determined based on acceptance by one of the classifier cascades.

10

The digital image may be one of multiple images in a series that include the face. The method may include correcting an uneven illumination condition of the face within a different image in the series than the digital image within which the illuminating condition is determined.

15

The uneven illumination classifier cascades may include a top illumination classifier, a bottom illumination classifier, and one or both of right and left illumination classifiers.

20

A further face detection method is provided. A digital image is acquired and a sub-window is extracted from the image. Separate sets of two or more shortened face detection classifier cascades are applied. One of the sets is trained to be selectively sensitive to a directional facial illumination, and another set of face detection classifier cascades is insensitive to directional facial illumination. The face data are identified as corresponding to a projection of the face within the digital image on one or a combination of the directional illumination classifier cascades plus a constant vector representing the face according to one or a combination of the directional illumination insensitive classifier cascades, thereby decomposing the face data into orthogonal subspaces for directional and uniform conditions. Based on the applying and identifying, a probability is determined that a face having a certain form of directional facial illumination is present within the sub-window. Based on the determining, an extended face detection classifier cascade trained for sensitivity to the form of directional facial illumination is applied. A final determination is provided that a face exists within the image sub-window. The process is repeated for one or more further sub-windows from the image or one or more further directional facial illuminations, or both.

30

The digital image may be one of multiple images in a series that include the face. The method may include correcting an uneven illumination condition of the face within a different image in the series than the digital image within which the illuminating condition is determined.

The directional illumination classifier cascades comprise a top illumination classifier, a bottom illumination classifier, and one or both of right and left illumination classifiers.

An illumination condition of a face within a sub-window may be determined based on acceptance by one or a combination of the classifier cascades.

A digital image acquisition device is provided with an optoelectronic system for acquiring a digital image, and a digital memory having stored therein processor-readable code for programming the processor to perform any of the methods described herein.

BRIEF DESCRIPTION OF THE DRAWINGS

Figures 1 and 2 illustrate different individuals at constant frontal illumination based on which a shape model and a texture model may be built.

Figure 3 illustrates images with various directional lighting conditions.

Figure 4 illustrates obtaining a residual texture g_{res} based on a projection of texture vectors on a uniform lighting (or illumination) subspace.

Figure 5A illustrates fitting a model in accordance with certain embodiments on a new image in which a spotlight is present at a person's left side.

Figure 5B illustrates a face patch that is correctly segmented, but wherein the person's identity is not so accurately reproduced in the synthesized image.

Figure 5C illustrates the extraction of real texture of an original image from inside a fitted

shape.

Figure 5D illustrates results for setting the parameters to zero and obtaining illumination normalization.

5

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Techniques are provided below wherein a texture space is decomposed into two orthogonal subspaces, one of uniform illumination and the second for illumination variability. An advantage of this approach is that two separate sets of parameters are used to control variations between individuals and variations in illumination conditions. Another advantage is that an exhaustive image database for training the statistical model is not needed. Statistical appearance models are described briefly, as well as a method for separating directional lighting variation. It is also described how to fuse two separate texture models. Examples of how the model may be fit are also described below and some experimental results are presented.

A linear texture model is constructed wherein the model is decomposed into orthogonal subspaces, one of which describes the model variability to directional changes in lighting conditions. By employing such a model various applications of linear texture models are made robust to variations in directional lighting conditions. Where such models are integrated within an image acquisition appliance such as a digital camera, external processing device, printer, display or other rendering appliance, or computer or other processing appliance, the models may be employed to improve image quality at the time of acquisition.

25

A linear texture model, specifically an active appearance model (AAM) is trained in a manner which separates the directional lighting variability from a second aspect of the model. The second aspect of the model may now be determined independently of the directional lighting applied to a scene. When incorporated within a digital camera such a model is substantially more robust to lighting variations and provides more acceptable image improvements.

30

As this describes a technique which can remove "directional lighting variability" from faces,

it can be combined with any of a wide various face analysis techniques; such as those described in US published applications nos. 2005/0140801, 2005/0041121, 2006/0204055, 2006/0204110, PCT/US2006/021393, 2006/0120599, 2007/0116379, PCT/US2007/075136, 2007/0116380, 2007/0189748, 2007/0201724, 2007/0110305, and 2007/0160307, and US applications nos. 11/761,647, 11/462,035, 11/767,412, 11/624,683, 11/753,397, 11/766,674, 11/773,815, 11/773,855, 60/915,669, 60/945,558, and 60/892,881, and US patents nos. 7,315,631, 7,336,821, 6,407,777 and 7,042,511.

Techniques are described herein that have reduced complexity and that serve to decompose a linear texture space of a facial appearance model into two (or more) linear subspaces, for example, one for inter-individual variability and another for variations caused by directional changes of the lighting conditions (others may include pose variation and other variables described herein). The approach used is to create one linear subspace from individuals with uniform illumination conditions and then filter a set of images with various directional lighting conditions by projecting corresponding textures on the previously built space. The residues are further used to build a second subspace for directional lighting. The resultant subspaces are orthogonal, so the overall texture model can be obtained by a relatively low complexity concatenation of the two subspaces. An advantage of this representation is that two sets of parameters are used to control inter-individual variation and separately intra-individual variation due to changes in illumination conditions.

STATISTICAL FACE MODELS OF SHAPE AND TEXTURE

Shapes are defined as a number of landmarks used to best describe the contour of the object of interest (i.e. the face). A shape vector is given by the concatenated coordinates of all landmark points, as $(x_1, x_2, \dots, x_L, y_1, y_2, \dots, y_L)^T$, where L is the number of landmark points.

The shape model is obtained by applying PCA on the set of aligned shapes

$$\mathbf{s} = \underline{\mathbf{s}} + \mathbf{PHI}_S \mathbf{b}_S \quad (1)$$

where $\underline{\mathbf{s}}$ is the mean shape vector, with N_S the number of shape observations; \mathbf{PHI}_S is the matrix having as its columns the eigenvectors; \mathbf{b}_S defines the set of parameters of the shape

model.

Texture, defined as pixel values across an object of interest, may also be statistically modelled. Face patches are first warped into the mean shape based on a triangulation algorithm. Then, a texture vector $(t_1, t_2, \dots, t_P)^T$ is built for each training image by sampling the values across the warped (shape-normalized) patches, with P being the number of texture samples.

The texture model is also derived by means of PCA on the texture vectors:

$$\mathbf{t} = \underline{\mathbf{t}} + \mathbf{PHI}_t \mathbf{b}_t \quad (2)$$

where $\underline{\mathbf{t}}$ is the mean texture vector, N_t being the number of texture observations; \mathbf{PHI}_t is again the matrix of eigenvectors, and \mathbf{b}_t the parameters for the texture model.

Sets of shape and texture parameters $\mathbf{c} = (\mathbf{W}_s \mathbf{b}_s; \mathbf{b}_t)$ are used to describe the overall appearance variability of the modelled object, where \mathbf{W}_s is a vector of weights used to compensate for the differences in units between shape and texture parameters.

TEXTURE SPACE DECOMPOSITION

To construct texture models, images from the Yale Face Database B were used (see below, and A.S. Georgiades, P.N. Belhumeur, and D.J. Kriegman, "From few to many: Generative models for recognition under variable pose and illumination," in FG, 2000, pp. 277–284). First, a shape model and a texture model are built from different individuals at constant frontal illumination (see Figure 1 and Figure 2). The shape and texture models are built as described in formulas (1) and (2), keeping also the same notations. This type of illumination stands as an approximation for uniform illumination conditions. The resultant texture eigenspace is referred to as the Uniform Lighting Subspace (ULS), or uniform illumination subspace, of the individuals.

For each individual, images are now considered with various directional lighting conditions (see Figure 3). The same reference shape is used to obtain the new texture vectors \mathbf{g} , which

ensures that the previous and new texture vectors have equal lengths.

These vectors are then filtered by projecting them on ULS (see formulas (3) and (4), below):

$$5 \quad \mathbf{b}_{t_opt} = \mathbf{PHI}_t^T (\mathbf{g} - \mathbf{t}) \quad (3)$$

$$\mathbf{g}_{filt} = \mathbf{t} + \mathbf{PHI}_t \mathbf{b}_{t_opt} \quad (4)$$

The residual texture is given by formula (5) below:

$$10 \quad \mathbf{g}_{res} = \mathbf{g} - \mathbf{g}_{filt} = \mathbf{g} - \mathbf{t} - \mathbf{PHI}_t \mathbf{b}_{t_opt} \quad (5)$$

The residues are used to create an orthogonal texture subspace of directional lighting. This may be referred to as a Directional Lighting Subspace (DLS) or directional illumination subspace. The directional lighting texture model is described, similar to formula (2) above, by formula (6) below:

$$\mathbf{g}_{res} = \mathbf{g}_{res} + \mathbf{PHI}_g \mathbf{b}_g \quad (6)$$

FUSION OF TEXTURE MODELS

20 DLS is built from the residual (difference) images subsequent to a projection on ULS. Thus, DLS is orthogonal to ULS. The fused texture model is given by formula (7) below:

$$\mathbf{t}_{fused} = \mathbf{t} + \mathbf{PHI}_t \mathbf{b}_t + \mathbf{PHI}_g \mathbf{b}_g \quad (7)$$

25

The fusion between the two texture models may be realized by, for example, a weighted concatenation of parameters. A vector of weighted shape parameters concatenated with the texture parameters is as follows in formula (8):

$$30 \quad \mathbf{c} = (\mathbf{W}_s \mathbf{b}_s; \mathbf{b}_t; \mathbf{W}_g \mathbf{b}_g) \quad (8)$$

where \mathbf{W}_s and \mathbf{W}_g are two vectors of weights used to compensate for the differences in units between the two sets of texture parameters, and for the differences in units between shape and texture parameters, respectively.

5

MODEL FITTING AND RESULTS

The Active Appearance Model (AAM) is a common technique used to optimize parameters of a statistical model of appearance as in G.J. Edwards, C.J. Taylor, and T.F. Cootes, "Interpreting face images using active appearance models," in *Proc. of the 3rd Int. Conf. on Automatic Face and Gesture Recognition*, Nara, Japan. As demonstrated by A.U. Batur and
10 M.H. Hayes, "Adaptive active appearance models," vol. 14, no. 11, pp. 1707–1721, November 2005, the standard AAM algorithm, which uses a gradient estimate built from training images, cannot generally be successfully applied on images when important variations in the illumination conditions are present. This is because the estimated gradient specializes around the mean of the dataset it is built from.

15

The solution proposed by Batur et al. is based on using an adaptive gradient AAM. The gradient matrix is linearly adapted according to texture composition of the target image, in order to generate a better estimate of the actual gradient. This technique represents a trade-off between using a fixed gradient (AAM) and numerically computing the gradient matrix at
20 each iteration (the standard optimization technique).

YALE FACE DATABASE B

To build exemplary statistical models, the standard Yale Face Database B is used and
25 referred to in the examples that follow. The database contains 5760 single light source images of 10 subjects each seen under 576 viewing conditions (9 poses x 64 illumination conditions). The images in the database were captured using a purpose-built illumination rig. This rig is fitted with 64 computer controlled strobes. The 64 images of a subject in a particular pose were acquired at camera frame rate (30 frames/second) in about 2 seconds, so there is only
30 small change in head pose and facial expression for those 64 (+1 ambient) images. The acquired images are 8-bit (gray scale) and the size of each image is 640(w) x 480(h).

Being a fairly comprehensive face database from the point of view of the range of directional lighting variations, it may be used for generating a steady directional lighting model. For this purpose, the frontal pose images are generally selected for use under all captured illumination conditions. This model can be further incorporated into a full (lighting-enhanced) face model.

EXEMPLARY FITTING SCHEME

A modified AAM-based fitting scheme is adapted for the lighting-enhanced face model described above. Having obtained a model in which there is a clear separation between identity and lighting variation facilitates the design of an AAM-based fitting scheme which could successfully be applied with this particular type of model.

Evaluation is performed of the lighting variations which are present inside a raw texture vector, based on the specialized illumination subspace/model. The illumination model is kept fixed throughout designing and training subsequent models for other types of variations (e.g. an enhanced identity model). This provides that the orthogonality condition and consistency of reference shape used for generating texture vectors are maintained.

Based on a direct projection of the original texture vector onto the illumination subspace, the lighting variations present in the original texture vector are estimated. The nature and extent of lighting variations is encoded in the amplitude of the parameters of the illumination model. The accuracy of this estimation is dependent on the comprehensiveness of the illumination model.

As a separate application, after estimating the lighting parameters, they can be reversed and thus a new texture vector can be generated which theoretically has zero lighting variations (in practice these lighting variations will only be removed to a certain extent).

The standard AAM technique is used now to drive all model parameters other than the lighting parameters. The lighting parameters (obtained by direct projection) are further employed for generating the complete synthetic texture. The modelled texture vector has now incorporated also the lighting variations, being thus adapted to the specific lighting conditions

present inside the original image. The resulting error surface is thus compatible with the constant gradient assumption of the AAM fitting technique. This is because the important directional lighting variations, which were practically responsible for the failure of a standard AAM fitting technique, are no longer reflected in the error vector. This step is introduced into the iterative optimisation scheme of a standard AAM.

It was experimentally confirmed that it suffices to have a low dimensional (possibly as low as 3-dimensions) lighting subspace for the proposed scheme to work. The low dimensionality of the lighting subspace ensures that the added computational costs involved (due to the multiple projections) do not affect significantly the overall fitting rate of the face model.

Enhancing the identity model can be realised by creating a more extensive training dataset. In order to insure that the orthogonality condition is still maintained, each new image is first projected on the lighting texture subspace in order to filter out the unwanted variation. The difference images are this time used to re-estimate identity subspace. For that, an iterative algorithm is employed to obtain the new mean of the overall texture model.

CLOSE INITIALIZATION REQUIREMENT

One weakness of the fitting technique proposed above is the requirement for a more accurate initialization of the model inside the image frame; namely the position and size of the face. The practical solution proposed here is to firstly employ the face detector described in P. A. Viola and M. J. Jones, "Robust real-time face detection." International Journal of Computer Vision, vol. 57, no. 2, pp. 137–154, 2004, and then tune its returned estimates of these face localization parameters with the centre of gravity and size of the specific reference shape used with the face model. The Viola-Jones face detector, based on the AdaBoost algorithm described, for example, in Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in European Conference on Computational Learning Theory, 1995, pp. 23–37, is thus firstly applied to estimate the centre and size of the face. The initialization step of the statistical model inside the image frame is tuned with the face detector using a statistical relationship learned from the training set. A statistical relation between the face detector estimates for the face position and size (rectangle region) and the position and size of the reference shape inside the image frame is

initially learnt (offline) from a set of training images. This relation is then used to obtain a more accurate initialization for the reference shape, tuned with the employed face detection algorithm. It is also important to have a reasonably close initialization to the real values in order to insure the convergence of the fitting algorithm.

5

EXTENSION OF THE LIGHTING-ENHANCED MODEL FOR COLOUR IMAGES

A notable applicability of the directional lighting sub-model (generated from a greyscale database) is that it can effectively be incorporated into a colour face model. The justification is that, theoretically, directional (white) lighting variations produce changes in the intensity component of a colour image, which is thus a problem that could be dealt with on greyscale. Yet, it is found in practice that, when working with colour images, directional lighting variations affect not only the intensity component but produce similar, although usually less significant effects on the chrominance components as well. Nonetheless, a colour face model designed principally based on the same technique described above for the greyscale, but employing instead colour texture vectors, is proven to offer significant benefits in terms of segmentation accuracy over the standard AAM-based equivalent.

Although more computationally expensive, the standard optimization technique may be used as it also provides enhanced accuracy. A fast optimization algorithm may also be used, wherein the AAM technique is used for optimizing model parameters other than illumination parameters. The model convergence may be applied on a new set of images with some variations of the illumination conditions, although the inventors found in their research no very important differences in the quality of the fit between the fused model and a standard model based on a single texture eigenspace built from the combined set of training images.

An advantage of fitting a fused model is that it offers control on illumination, based on the separate set of illumination parameters. Thus, after fitting the model, one can obtain an estimate of the illumination conditions present, normalize the illumination (e.g. to uniform illumination), or generate different illumination conditions.

In Figures 5A-5D, an example is illustrated where a model as described herein is fit on a new image in which a spotlight is present at the person's left side as in Figure 5A. As it can be

seen from Figure 5B, the face patch is correctly segmented, yet the person's identity is not so accurately reproduced in the synthesized image. This is due to the limitations of the ULS model which was built using a small number of observations. Yet one can now extract the real texture of the original image from inside the fitted shape as illustrated at Figure 5C. The real texture vector can be viewed as the projection of the individual on the directional lighting eigenspace plus a constant vector representing the individual under uniform lighting. This can be written as in formula (9) below:

$$\mathbf{t}_{\text{real}} = \mathbf{t}_{\text{unif}} + \mathbf{PHI}_g \mathbf{b}_{g_opt} \quad (9)$$

where \mathbf{b}_{g_opt} are the illumination parameters which were estimated during the overall optimization stage. By altering \mathbf{b}_{g_opt} new illuminations can be generated. Figure 5D shows the results for setting all parameters to zero and obtaining illumination normalization.

A statistical face model based on texture space decomposition as has been described enables the separation of illumination variability and intra-individual variability. The model is useful in applications where the control over the illumination conditions is desired. The separation of the sets of parameters is also useful when it is not desirable to maintain an exhaustive image database for all the modelled components. Thus, an appropriate (separate) database can be used for modelling components like shape (including pose) variability, inter-individuals variability, and/or directional lighting. An improved shape model, designed to include also pose variability, can also be used to further enhance the capabilities of the overall appearance model.

The technique can also be extended to color images using a color training database for modelling individual variability under uniform lighting. Advantageous models provided herein offer solutions in applications like face recognition and/or face tracking for dealing with variations in pose and/or illumination conditions, where many of the current techniques still show weak points.

ALTERNATIVE EMBODIMENTS

In-camera and external image processing applications may be combined with features provided herein in further embodiments. In an image processing embodiment, features of a linear texture or AAM model as described above or elsewhere herein are enhanced in combination with a program that accurately removes unnecessary periodic noise components from an image (see US published patent application no. 2006/0257047). A reconstruction unit may generate a reconstructed image without a periodic noise component by fitting to a face region detected in an image by a face detection unit a mathematical model generated according a method of AAM using a plurality of sample images representing human faces without a periodic noise component. The periodic noise component may be extracted by a difference between the face region and the reconstructed image, and a frequency of the noise component may be determined. The noise component of the determined frequency may then be removed from the image.

A further embodiment relating to photography and photography programs includes a program which determines an exposure value based on a face region with high accuracy and with less effect of a background region or density contrast caused by shadow (see US published patent application no. 2006/0268150). A face detection unit may be used to detect a face region from a face candidate region by fitting to the face candidate region a mathematical model generated by a method of AAM using a plurality of sample images representing human faces. An exposure value determination unit may then determine an exposure value for photography based on the face region.

A further embodiment relating to photography includes a program wherein an auxiliary light source such as a flash is controlled for highly accurately reflecting a lighting condition with less effect of a factor other than the lighting condition such as contrasting density caused by shadow or black compression or white compression (see US published patent application no. 2006/0269270). In this embodiment, a parameter acquisition unit obtains weighting parameters for principal components representing lighting conditions in a face region in an image detected by a face detection unit by fitting to the detected face region a mathematical model generated according to a method of AAM using a plurality of sample images representing human faces in different lighting conditions. A flash control unit may obtain a flash control condition by referring to a reference table according to the parameters, and may control flash light emission according to the control condition.

Another embodiment involves image processing and includes a program wherein sharpness is adjusted for enhanced representation of a predetermined structure in an image. A parameter acquisition unit may obtain a weighting parameter for a principal component representing a degree of sharpness in a face region detected by a face detection unit as an example of the predetermined structure in the image by fitting to the face region a mathematical model generated by a statistical method such as AAM based on a plurality of sample images representing human faces in different degrees of sharpness (see US published patent application no. 2007/0070440). Based on a value of the parameter, sharpness may be adjusted in at least a part of the image. For example, a parameter changing unit may change the value of the parameter to a preset optimal face sharpness value, and an image reconstruction unit may reconstruct the image based on the parameter having been changed and may output the image having been subjected to the sharpness adjustment processing.

A further embodiment involves removal of an unnecessary or undesirable blemish in a digital image, such as an acne blemish or a dust artefact, i.e., a blemish caused by the presence of dust in the air or on an optical surface between the object and the image sensor. The blemish may be removed completely from a predetermined structure such as a face in a photograph image without manual operation and skills (see US published patent applications 2005/0068452 and 2007/0071347). A blemish removal unit may be fit to a face region as the structure in the image. A mathematical model may be generated according to a statistical method such as AAM using sample images representing the structure without the component to be removed, and an image reconstruction unit may reconstruct an image of the face region based on parameters corresponding to the face region obtained by the fitting of the model. An image may be generated by replacing the face region with the reconstructed image. Since the mathematical model has been generated from the sample images of human faces without blemishes caused by dust or acne, the model does not include such blemishes. Therefore, the reconstructed face image generated by fitting the model to the face region does not include such blemishes.

In another image processing embodiment, features of a linear texture or AAM model as described above or elsewhere herein are enhanced in combination with a program that reduces graininess with accuracy by finding a degree of graininess in an image (US published

patent application no. 2006/0291739). A parameter acquisition unit may obtain a weighting parameter for a principal component representing a degree of graininess in a face region found in the image by a face detection unit by fitting to the face region a mathematical model generated by a method of AAM using a plurality of sample images representing human faces in different degrees of graininess. A parameter changing unit may change the parameter to have a desired value. A graininess reduction unit reduces graininess of the face region according to the parameter having been changed.

Another embodiment includes a combination of one or more features of a linear texture or AAM model as described above or elsewhere herein with a program that converts the resolution of an input image using a certain AAM method (see US published patent application no. 2006/0280380). A resolution conversion unit is described for converting a resolution of an image having been subjected to correction. A face detection unit is described for detecting a face region in the resolution-converted image. A reconstruction unit fits to the face region detected by the face detection unit a mathematical model generated through the method of AAM using sample images representing human faces having the same resolution as the image, and reconstructs an image representing the face region after the fitting. In this manner, an image whose resolution has been converted is obtained.

In addition the invention may be applied as a pre-processing step to a face recognition system. In such an embodiment face regions are detected within an acquired image using techniques such as may be described in US patent no. 7,315,631. One or more AAM models are also applied to the determined face region when these techniques are used.

In one embodiment the model components which are members of the directional lighting dependent subspace are simply discarded leaving the lighting independent components. Face recognition is then performed on these filtered components. Naturally the face recognizer is initially trained on a database of similarly filtered images. This technique may be applied to most known face recognition techniques including, DCT, PCA, GMM and HMM based face recognizers.

In an alternative embodiment, the directional lighting dependent parameters are adjusted to a set of mean values and a modified face region is created. This modified face region is then

passed to a face recognition engine.

A corresponding separation of pose variation into an orthogonal subspace is also achieved in accordance with another embodiment. This enables a similar separation of facial pose and other model characteristics. Thus a corresponding set of AAM models which describe variations in some characteristic(s) of a face region in a pose-independent manner are provided in this embodiment. These models can be combined with the illumination independent models, as described above for example.

In one embodiment, a first AAM model is applied wherein one or more pose-dependent parameters are adjusted to a set of mean values and a modified face region is created. A second AAM model is applied wherein one or more directional lighting dependent parameters are adjusted to a set of mean values and a second modified face region is created. This modified face region is then provided to a frontal face recognizer.

While an exemplary drawings and specific embodiments of the present invention have been described and illustrated, it is to be understood that that the scope of the present invention is not to be limited to the particular embodiments discussed. Thus, the embodiments shall be regarded as illustrative rather than restrictive, and it should be understood that variations may be made in those embodiments by workers skilled in the arts without departing from the scope of the present invention as set forth in the claims that follow and their structural and functional equivalents.

In addition, in methods that may be performed according to the claims below and/or preferred embodiments herein, the operations have been described in selected typographical sequences. However, the sequences have been selected and so ordered for typographical convenience and are not intended to imply any particular order for performing the operations, unless a particular ordering is expressly provided or understood by those skilled in the art as being necessary.

Claims:

1. A method of determining a characteristic of a face or certain other object within a scene captured in a digital image, comprising:
 - 5 (a) acquiring a digital image including a face or certain other object within a scene;
 - (b) applying a linear texture model that is constructed based on a training data set and that comprises a class of objects including a first subset of model components that exhibit a dependency on directional lighting variations and a second subset of model components which are independent of directional lighting variations;
 - 10 (c) determining an initial location of the face or certain other object in the scene;
 - (d) obtaining a fit of said model to said face or certain other object including adjusting one or more individual values of one or more of the model components of said linear texture model;
 - (e) based on the obtained fit of the model to said face or certain other object in the
15 scene, determining at least one characteristic of the face or certain other object; and
 - (f) electronically storing, transmitting, applying a face or other object recognition program to, editing, or displaying the corrected face image or certain other object including the determined characteristic, or combinations thereof.
- 20 2. The method of claim 1, wherein the model components comprise eigenvectors, and the individual values comprises eigenvalues of the eigenvectors.
3. The method of claim 1, wherein the at least one determined characteristic comprises a
25 feature that is independent of directional lighting.
4. The method of claim 1, further comprising generating a reconstructed image without a periodic noise component, including obtaining a second fit of the face or certain other object to a second linear texture model that is based on a training data set and that comprises a class
30 of objects including a set of model components which are without a periodic noise component.
5. The method of claim 4, further comprising:
 - (i) extracting the periodic noise component including determining a difference

between the face or certain other object and the reconstructed image, and

(ii) determining a frequency of the noise component; and

(iii) removing the periodic noise component of the determined frequency.

5 6. The method of claim 1, further comprising determining an exposure value for the face or certain other object, including obtaining a fit to the face or certain other object to a second linear texture model that is based on a training data set and that comprises a class of objects including a set of model components that exhibit a dependency on exposure value variations.

10 7. The method of claim 6, further comprising including reducing an effect of a background region or density contrast caused by shadow or both.

8 The method of claim 1, further comprising controlling a flash to accurately reflect a lighting condition, including obtaining a flash control condition by referring to a reference
15 table and controlling a flash light emission according to the flash control condition.

9. The method of claim 8, further comprising reducing an effect of contrasting density caused by shadow or black compression or white compression or combinations thereof

20 10. The method of claim 1, further comprising adjusting or determining a sharpness value, or both, including:

(i) applying a second linear texture model that is constructed based on a training data set and comprises a class of objects including a subset of model components that exhibit a dependency on sharpness variations,

25 (ii) obtaining a fit of said second model to said face or certain other object in the scene including adjusting one or more individual values of one or more model components of said second linear texture model; and

(e) based on the obtained fit of the second model to said face or certain other object in the scene, adjusting a sharpness of the face or certain other object including changing one or
30 more values of one or more model components of the second linear texture model to generate a further adjusted object model.

11. The method of claim 1, further comprising removing a blemish from a face or certain

other object, including:

(i) applying a second linear texture model that is constructed based on a training data set and comprises a class of objects including a subset of model components that do not include such a blemish;

5 (ii) obtaining a fit of said second model to said face or certain other object in the scene including adjusting one or more individual values of one or more model components of said second linear texture model; and

(e) based on the obtained fit of the second model to said face or certain other object in the scene, removing the blemish from the face or certain other object including changing one
10 or more values of one or more model components of the second linear texture model to generate a further adjusted object model.

12. The method of claim 11, wherein the blemish comprises an acne blemish or other skin blemish.

15

13. The method of claim 11, wherein the blemish comprises a photographic artefact.

14. The method of claim 1, further comprising adjusting or determining a graininess value, or both, including:

20 (i) applying a second linear texture model that is constructed based on a training data set and comprises a class of objects including a subset of model components that exhibit a dependency on graininess variations,

(ii) obtaining a fit of said second model to said face or certain other object in the scene including adjusting one or more individual values of one or more model components of said
25 second linear texture model; and

(e) based on the obtained fit of the second model to said face or certain other object in the scene, adjusting a graininess of the face or certain other object including changing one or more values of one or more model components of the second linear texture model to generate a further adjusted object model.

30

15. The method of claim 1, further comprising converting, adjusting or determining a resolution value, or combinations thereof, including:

(i) applying a second linear texture model that is constructed based on a training data

set and comprises a class of objects including a subset of model components that exhibit approximately a same resolution as said face or certain other object,

(ii) obtaining a fit of said second model to said face or certain other object in the scene including adjusting one or more individual values of one or more model components of said

5 second linear texture model; and

(e) based on the obtained fit of the second model to said face or certain other object in the scene, converting a resolution of the face or certain other object including changing one or more values of one or more model components of the second linear texture model to generate a further adjusted object model.

10

16. A digital image acquisition device including an optoelectronic system for acquiring a digital image, and a digital memory having stored therein processor-readable code for programming the processor to perform a method of determining a characteristic of a face or certain other object within a scene captured in a digital image, according to any of the above

15 claims.




Figure 1




Figure 2

2 / 3

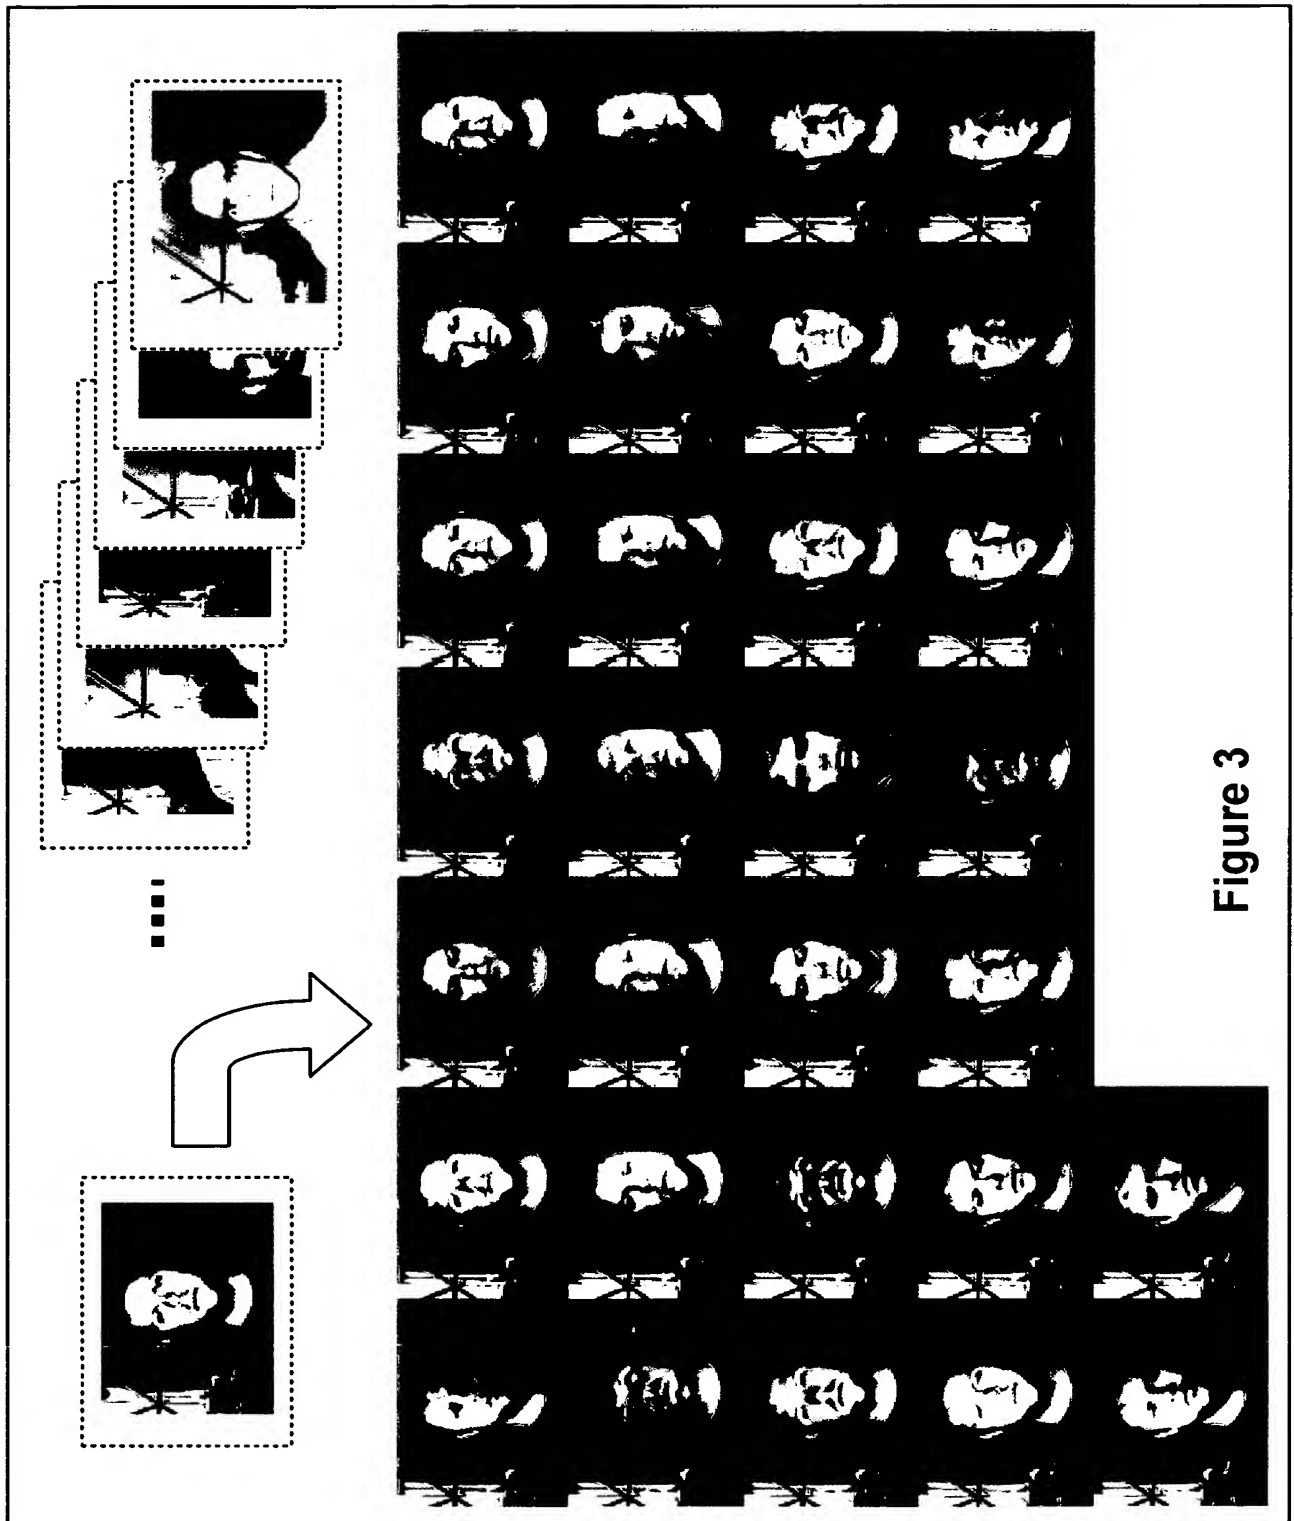


Figure 3

3 / 3

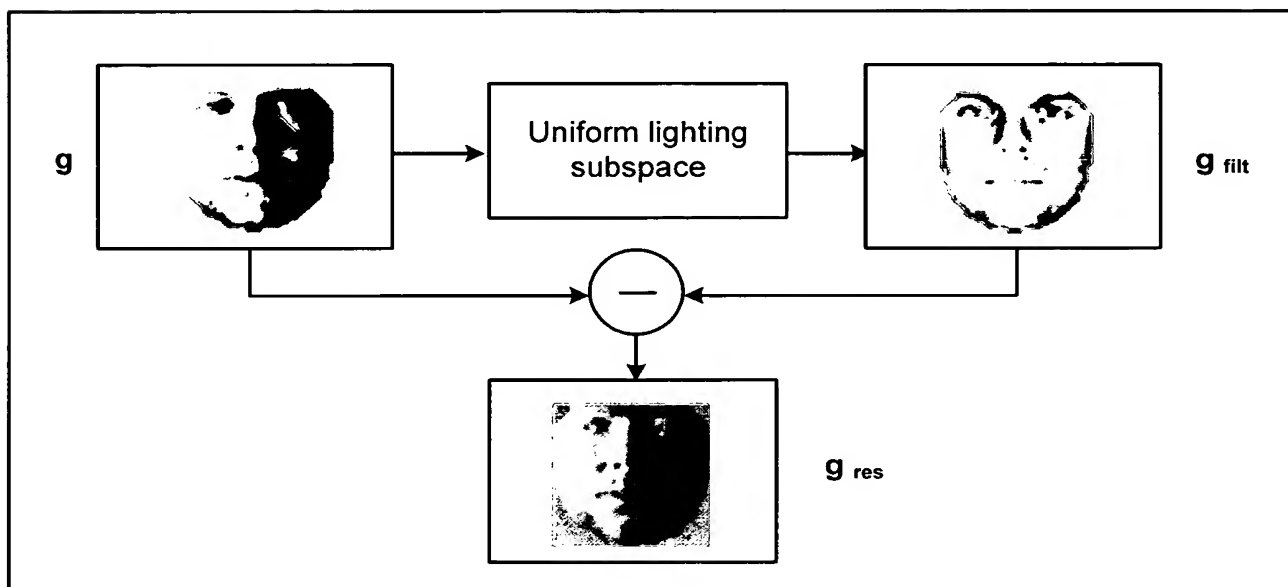


Figure 4



Figure 5A



Figure 5B




Figure 5C



Figure 5D